arXiv:1603.03040v1 [astro-ph.SR] 9 Mar 2016

The Cannon 2: A data-driven model of stellar spectra for detailed chemical abundance analyses

Andrew R. Casey[1], David W. Hogg[2,3,4,5], Melissa Ness[5],
Hans-Walter Rix[5], Anna Y. Q. Ho[6], and Gerry Gilmore[1]

arc@ast.cam.ac.uk

ABSTRACT

We have shown that data-driven models are effective for inferring physical attributes of stars (labels; $T_{\rm eff}$, $\log g$, [M/H]) from spectra, even when the signal-to-noise ratio is low. Here we explore whether this is possible when the dimensionality of the label space is large ($T_{\rm eff}$, $\log g$, and 15 abundances: C, N, O, Na, Mg, Al, Si, S, K, Ca, Ti, V, Mn, Fe, Ni) and the model is non-linear in its response to abundance and parameter changes. We adopt ideas from compressed sensing to limit overall model complexity while retaining model freedom. The model is trained with a set of 12,681 red giant stars with high signal-to-noise spectroscopic observations and stellar parameters and abundances taken from the *APOGEE* Survey. We find that we can successfully train and use a model with 17 stellar labels. Validation shows that the model does a good job of inferring all 17 labels (typical abundance precision is 0.04 dex), even when we degrade the signal-to-noise by discarding $\gtrsim$50% of the observing time. The model dependencies make sense: the spectral derivatives with respect to abundances correlate with known atomic lines, and we identify elements belonging to atomic lines that were previously unknown. We recover (anti-)correlations in abundance labels for globular cluster stars, consistent with the literature. However we find that the intrinsic spread in globular cluster abundances is 3–4 times smaller than previously

[1]Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

[2]Simons Center for Data Analysis, 160 Fifth Avenue, 7th Floor, New York, NY 10010, USA

[3]Center for Cosmology and Particle Physics, Department of Physics, New York University, 4 Washington Pl., room 424, New York, NY, 10003, USA

[4]Center for Data Science, New York University, 726 Broadway, 7th Floor, New York, NY 10003, USA

[5]Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany

[6]Astronomy Department, California Institute of Technology, MC 249-17, 1200 East California Blvd, Pasadena, CA 91125, USA

reported. We deliver 17 labels with associated errors for 87,563 red giant stars, as well as open-source code to extend this work to other spectroscopic surveys.

1. Introduction

The detailed chemical composition of a star's photosphere reflects its formation environment, and the mix of supernovae that preceded it. This photospheric *chemical fingerprint* remains largely unchanged throughout a star's lifetime, providing a fossil record of local star-formation history. Although the differences in detailed chemical abundances may be subtle between two formation sites, if stars could be linked to their natal gas cloud by their chemical fingerprint, a sufficiently large collection of precise stellar abundances would unravel the complete chemical evolution of the Milky Way.

This goal has only recently become feasible with the increasing volume and quality of stellar spectra obtained in the last decade. Specifically, large surveys are obtaining high-resolution ($\mathcal{R} \gtrsim 20,000$), high signal-to-noise (S/N) ratio spectra for $\sim 10^5$—10^6 stars across all components of the Galaxy (Gilmore et al. 2012; Zasowski et al. 2013; De Silva et al. 2015). This is a sharp relative increase in data volume; it has mandated the automation of spectral analysis, and encouraged dozens of groups to produce bespoke pipelines.

Most automated pipelines have grown from classical, manual methods. There has been relatively little work on unconventional methods to analyse spectra. Spectroscopists have instead sought to code their experience (or subjectivity), with heuristics enforced for wavelength masks, convergence criteria, and similar analysis issues. Most of these decisions are based on the optimization of accuracy for particular stars with good quality data (e.g., Solar-like stars with high S/N ratios). As a consequence, these heuristics are frequently incompatible for data with more modest (and representative) S/N ratios. Indeed, it can be shown from repeat or blind experiments that traditional pipelines routinely yield imprecise abundances for noisy data. Moreover, the results from individual pipelines are inconsistent, with differences an order of magnitude larger than what state-of-the-art spectroscopic studies seek to measure (e.g., the effects of atomic diffusion, evolutionary mixing, planetary accretion). Thus, while there has been substantial effort to automate classical analysis techniques, they are generally imprecise at modest S/N ratios, and frequently yield inaccurate (biased) results.

Considerable effort has been spent on improving the accuracy of physical models of stars. Indeed it is impossible to measure physical properties of stars (or their chemical abundances) accurately without accurate physical models of stellar spectra. However physics-based mod-

els of stars have a number of known problems: they rely on incorrect approximations (e.g. one-dimensional atmospheres) and incomplete laboratory data. Three-dimensional models remain computationally impractical for more than just a few stars. As a consequence of the limited atmosphere dimensionality, crude (knowingly incorrect) approximations for the convection are necessary. Although some grids of three-dimensional hydrodynamic models have been produced and averaged to one-dimensional approximations, these models assume local thermodynamic equilibrium (LTE). Properly accounting for departures from LTE is a formidable analytic and computational challenge. Additionally, while laboratory efforts have thoroughly improved much of the faulty atomic and molecular data, this process is unquestionably incomplete. For these reasons there are some spectral features that are much better understood than others. As a consequence, physics-based methods are restricted to spectral regions and parameter spaces that are understood marginally better, which limits their applicability and interpretability *by construction*.

In detail, physics-based models do not explain all pixels of a stellar spectrum at the precision with which we are currently observing. The data quality have outgrown the classical methods used to analyse them. This led to the creation of *The Cannon*[1] (Ness et al. 2015a), a data-driven—as opposed to physics-based—model for stellar spectra.

Before we continue to introduce background on this work, we first need to introduce some relevant terminology. Throughout this work we will call stellar parameters ($T_{\rm eff}$ and $\log g$) and the full set of 15 chemical abundances collectively "labels". This unifies and collapses the phrase "stellar parameters and chemical abundances" to a word, and connects to relevant terminology for supervised methods in the machine learning and statistics literatures. *The Cannon* uses spectra from stars that have labels known with high fidelity to train a model for stellar spectra. The trained model can then be used on new data, and precisely estimate labels.

The Cannon is a data-driven model, but it differs from standard machine learning approaches because it contains an explicit noise model. This means that *The Cannon* can transfer labels from high S/N training set stars to low S/N test set stars; that is, the training set and the test set do not need to be statistically identical. This is related to the fact that *The Cannon* is an interpretable causal model; the internals of the model are the dependencies of the spectral expectation and variance on wavelength and physical parameters of the star, plus an explicit noise model. Given a representative set of stars with known labels of high-

[1]It is important (to us) to note that *The Cannon* is named not after a weapon but instead after Annie Jump Cannon, who was the first to correctly order stellar spectra in temperature order (and who did so by looking at the data, and without any use of physics-based models, see, e.g., Cannon & Pickering 1912).

fidelity, *The Cannon* provides a generative model for stellar spectra based on a non-linear combination of the labels.

While a physical model is required to construct an acceptable training set for *The Cannon* to use, *The Cannon* also offers opportunities to reconcile known problems with physics-based models. It provides labels for training set data at wavelengths where there is limited or missing atomic information, allowing for every pixel to contribute in measuring labels from noisy spectra. If there are a sufficient number of stars observed by two surveys, *The Cannon* can also be used to determine labels that are consistent across surveys and wavelengths (Ho et al. 2016). Moreover, as we will show in this work, the model internals have physical interpretations, allowing for the identification of previously unknown spectral lines. This is a first step towards using *The Cannon* to improve physics-based models.

Only a small number of labels were used in the first work with *The Cannon*: three in the original work, and four or five in later work (Ness et al. 2015a,b; Ho et al. 2016). Here we were guided by thoughts related to density estimation. To sample a K-dimensional label space well, the size of the training set should scale exponentially (or worse) with K. Subsequent experiments, however, did not bear this out. We found that we can transfer many labels from the training set to the test set, with training sets of (just) thousands of stars. The fundamental reason is that *The Cannon* is *not* a density estimator. It is more like an *interpolator*, which effectively finds stars in the training set that are close to the test star, and transfers labels, using a polynomial model as a smooth interpolation function.

The capacity to extend to a larger set of labels without significant computational detriment offers tantalizing opportunities. The most straightforward is that it provides a measurement of many elemental abundances. In doing so, however, it can be shown that a standard *Cannon* model yields noisy coefficients (spectral derivatives; see next Section) that are incompatible with expectations from physics: the training coefficients may optimize to produce coefficients of abundance labels with non-zero contributions at *all pixels*, however physically we know that spectral lines of most elements do not contribute at all wavelengths.

For this reason alone we know that the problem ought to be sparse. Here we exploit this knowledge to the fullest, using standard regularization methods to discover and enforce sparsity. We consider the *entire* 17-dimensional label space produced by the *APOGEE ASPCAP* pipeline. For our training set, we adopt the *ASPCAP* labels of the stars with the highest S/N spectra. We show by validation that we can transfer these labels to much lower S/N stars, with reduced precision but no strong biases. We then use the system to label all of the stars in the *APOGEE* DR12 data set. After validating our model, we confirm our abundance precision using tests of globular clusters.

2. Method

Before we outline our assumptions, we need to define the different data sets we will use, as well as terminology related to our method. Consider a stellar spectroscopic survey. Within the survey data is a set of stars observed at high S/N, where the labels for those stars are known with high fidelity. We call this sample the *reference set*, and all other stars are set aside into the *test set*. Note that there may be labels reported for stars in the *test set*, but the implicit assumption is that those labels are not known with high fidelity. We are going to construct (train) *The Cannon* using a subset of the reference set, allowing us to *predict* stellar fluxes and *measure* (test) labels for stars in the test set. It is important to note that we do not use every star in the reference set for training: here we will define the *training set* as a random subset of the reference set. All stars in the reference set that are not part of the training set will form the *validation set*, which we will use to *validate* the predictive power of our model.

We assume the following about *The Cannon* and the *APOGEE* DR12:

- Stars with similar labels ($T_{\rm eff}$, $\log g$, and abundances) have similar spectra.
- The expectation that the spectrum of a star is a smooth function of the values of the labels for that star. Further than this, we assume that the function is so smooth it can be reasonably approximated with a quadratic form in label space.
- The resolution of all *APOGEE* spectra are identical and all spectra are calibrated to the same rest wavelength grid.
- The *APOGEE* noise variances reported are (nearly) correct, normally distributed, and independent from pixel to pixel. Importantly, we are *not* assuming that different stars have similar noise variances, nor that the reference and test sets have the same noise model.
- We have a training set of stars with *mean* accurate labels, where "accurate" is defined by the accuracy requirements of the output labels. It might be more appropriate to say that we are assuming that the training set stars have *consistent* labels (consistent with the assumptions of smoothness, above). Indeed, because *The Cannon* is a data-driven model we stress that it cannot create "ground truth": we are propagating labels from the training set.
- The training set is representative, in the sense that the training set stars span the label space similarly to how the test set spans the label space.

- We assume that our continuum normalization procedure (described below) is consistent. We do not require "true" continuum-normalization in the classical sense because any offset (even a label-dependent residual due to a strong absorption line) can be captured by the model. Instead we require that our normalization procedure is invariant with respect to S/N.

Given these assumptions, the model we adopt is

$$y_{jn} \quad = \quad \boldsymbol{v}(\ell_n) \cdot \boldsymbol{\theta}_j + e_{jn} \quad , \tag{1}$$

where y_{jn} is the data for star n at wavelength pixel j, $\boldsymbol{v}(\ell_n)$ is a function that takes as input the label list ℓ_n of length K for star n and outputs a vector of length $D > K$ of functions of those labels, $\boldsymbol{\theta}_j$ is a vector of length D of parameters controlling the model at wavelength pixel j, and e_{jn} is a noise draw or residual. We refer to $\boldsymbol{v}(\ell_n)$ as "the vectorizing function", which allows for arbitrarily complex functions that might not be simple polynomial expansions of the label list ℓ_n (e.g., sums of sines and cosines). Inasmuch as the model is good, the noise values e_{jn} can be taken to be drawn from a Gaussian with zero mean and variance $\sigma_{jn}^2 + s_j^2$, where σ_{jn}^2 is the pipeline-reported uncertainty variance on datum y_{jn} and s_j^2 is a parameter describing excess variance at wavelength pixel j.

Two comments about the model (1). The first is that, because the e_{jn} are thought of as being drawn from a probability density function (pdf), it is a probabilistic model for the spectral data y_{jn}. The second is that the output of the function $\boldsymbol{v}(\ell)$ can be thought of as a row of the "design matrix" that defines the possible freedom given to the spectrum expectation model.

In the *training step*, we fix the K-lists of labels ℓ_n for all training set stars n. We seek, at each wavelength pixel j, the $[D+1]$ parameters $\boldsymbol{\theta}_j, s_j^2$ that optimize a penalized likelihood:

$$\boldsymbol{\theta}_j, s_j^2 \quad \leftarrow \quad \underset{\boldsymbol{\theta}, s}{\mathrm{argmin}} \left[\sum_{n=0}^{N-1} \frac{[y_{jn} - \boldsymbol{v}(\ell_n) \cdot \boldsymbol{\theta}]^2}{\sigma_{jn}^2 + s^2} + \sum_{n=0}^{N-1} \ln(\sigma_{jn}^2 + s^2) + \Lambda_j \, Q(\boldsymbol{\theta}) \right] \quad , \tag{2}$$

where Λ_j is a regularization parameter, and $Q(\boldsymbol{\theta})$ is a regularizing function that encourages parameters to take on zero values. The regularizing function takes a D-vector as input and returns a scalar value. We call this penalized likelihood—the argument of the argmin in equation (2)—the *training scalar*. We will adopt for the regularizing function $Q(\boldsymbol{\theta})$ in the training scalar a modification of L1 regularization, discussed below. Although the training-step optimization problem will not in general be convex, we can make choices for $Q(\boldsymbol{\theta})$ (and we will) to make the problem such that it would be convex at any fixed value of s^2; for this reason it will tend to optimize well in most cases of interest. Convexity is

extremely advantageous: it ensures that only a single global minima exists, implying that any reasonable optimization routine is guaranteed to optimize to the correct global minima (i.e., no local minima).

The regularization parameter Λ_j sets the strength of the regularization; as Λ_j increases, the number of non-zero components of the parameter vector $\boldsymbol{\theta}_j$ will decrease. We give the regularization parameter a subscript j because in general we can set it differently at every wavelength. This makes sense, because different wavelengths have very different dependences on components of the label list ℓ. In practice the value of Λ_j should be set by full cross-validation, and possibly include restrictions based on physical arguments (e.g., a particular element does not have any spectral lines near this pixel, therefore the contributions from this label must be zero). This vastly expands the number of potential hyper-parameters, and thus the computing expense required to determine them. Thus for the purpose of this work we will set a single value of Λ (for all j pixels) by validation.

In the *test step*, we fix the parameters $\boldsymbol{\theta}_j, s_j^2$ at all wavelength pixels j. We seek, for each test set star m, the K-list of labels ℓ_m that optimizes the likelihood:

$$\ell_m \quad \leftarrow \quad \underset{\ell}{\mathrm{argmin}} \left[\sum_{j=0}^{J-1} \frac{[y_{jm} - \boldsymbol{v}(\ell) \cdot \boldsymbol{\theta}_j]^2}{\sigma_{jm}^2 + s_j^2} \right] \quad . \tag{3}$$

If the vectorizing function $\boldsymbol{v}(\ell)$ is non-linear (as it is in our quadratic model), the test-step optimization is not convex. However, because there are many pixels j acting, each of which has a different functional dependence on the labels in the label list ℓ, in practice the optimization finds a good value for the label list ℓ. We call this partial log likelihood—the argument of the argmin in equation (3)—the *test scalar*.

The model freedom of *The Cannon* is set by the vectorizing function $\boldsymbol{v}(\ell)$—which takes the K-element label list ℓ and expands it into a D-dimensional vector of components for the linear model—and the regularization $\Lambda_j\, Q(\boldsymbol{\theta})$. Because we want the (simple, see below) regularization to treat the different parameters (the different components of $\boldsymbol{\theta}$) in some sense "equally", we have to make sensible choices in the vectorizing function $\boldsymbol{v}(\ell)$. One thing that the vectorizing function $\boldsymbol{v}(\ell)$ can do is offset the labels by some kind of fiducial (mean, median, or other central) value, such that $\boldsymbol{v}(\ell) = 0$ is at a central location in the label space. Another is to divide out a scale, because, for example, $T_{\rm eff}$ values are in the thousands, but $\log g$ values are of order unity. If scale is not divided out, the (isotropic in $\boldsymbol{\theta}$) regularization will be much more harsh, effectively, on some parameters than others. An extension of this work might be to consider different regularization terms for each j pixel *and* every label in ℓ.

In what follows, we adopt the median value in the training set for each label value as

the fiducial offset. We choose a dimensionless scale factor f times the label range (defined as the difference between the 97.5th percentile and the 2.5th percentile of the training set along each label direction) as the scale such that, for example, the $T_{\rm eff}$ value for star n is rescaled as:

$$\hat{T}_{{\rm eff},n} = \frac{T_{{\rm eff},n} - p_{T_{\rm eff},50}}{f \cdot |p_{T_{\rm eff},97.5} - p_{T_{\rm eff},2.5}|} \quad (4)$$

where $p_{\theta,k}$ is the k-th percentile value of θ in the training set labels. For the regularizer $Q(\boldsymbol{\theta})$ we adopt a variant of L1 regularization; we set

$$Q(\boldsymbol{\theta}) \quad = \quad \sum_{d=1}^{D-1} |\theta_d| \quad , \quad (5)$$

where the sum is over the $[D-1]$ components of $\boldsymbol{\theta}$, excluding the zeroth component because we don't ever expect that component to vanish.[2] There is a great deal of theory about this kind of regularization; it is called L1 or the *LASSO* (Tibshirani 1996). L1 regularization encourages parameters to vanish precisely but doesn't break convexity for convex problems. It is important to note that the scale factor f plays into the regularization, because as the scale factor grows, the more penalized the cross-terms (e.g., $T_{\rm eff}$·[Al/H], [Mg/H]·[Si/H]) become relative to the linear terms. We will heuristically set the values of the hyper-parameters Λ and f by validation in Section 4.1.

3. Training, validation, and test data

We employ the *APOGEE* Data Release 12 data (Alam et al. 2015) to demonstrate the effectiveness of a regularized *Cannon* model in high-dimensional label space. We constructed the reference set using sensible criteria for stars analysed with version `v603` of the *ASPCAP* pipeline (Holtzman et al. 2015; García Pérez et al. 2015). We first removed any stars with problematic flags from the *ASPCAP* pipeline, i.e., we required `ASPCAPFLAG = 0`. We excluded stars with S/N ratios outside the range of 200-300, and stars with a radial velocity scatter

[2]Forgive a notational similarity here: When we subscript (bold) $\boldsymbol{\theta}$ with j we mean "the D-vector of parameters associated with wavelength pixel j"; when we subscript (non-bold) θ with d we mean "the single parameter along coordinate axis d in the D-dimensional $\boldsymbol{\theta}$ vector space."

larger than 1 km s^{-1}. We further demanded that our training set include reported abundances in all 15 elements (C, N, O, Na, Mg, Al, Si, S, K, Ca, Ti, V, Mn, Fe, and Ni), and restricted the abundance range such that $2 > [\mathrm{X/Fe}] > -2$, $[\mathrm{Fe/H}] > -3$, and $[\alpha/\mathrm{Fe}] > -0.1$. A visual comparison of the [V/H] labels with other (Fe-peak) abundance labels showed that many stars in the reference set had spurious measurements of [V/H]. For this reason we required that stars in the reference set have $[\mathrm{V/Fe}] > -0.6$. The distilled sample includes 14,141 red giant branch stars that will form our reference set, with [Fe/H] ranging from $[\mathrm{Fe/H}] = -2.10$ to $+0.30$ dex.

We randomly assigned each star in the reference set a uniformly-distributed integer q between 0 and 9, inclusive. We assign stars with $q > 0$ to the *training set*, and those with $q = 0$ to the *validation set*. We have ensured that the same q integer was assigned to each star between the different experiments described below. The training set (see Figure 1) comprises 12,681 red giant branch stars, and the validation set includes 1,460 stars.

Here we describe the requisite data processing steps before we can train the spectral model. The *APOGEE* `apStar` files contain rest-wavelength, resampled fluxes from individual visits for a given star, and an uncertainty array for those fluxes. Although the `aspcapStar` files contain stacked, pseudo-continuum normalized spectra for a given star, we chose to only use the fluxes from the `apStar` data files throughout this work. Our reasoning is as follows: The *ASPCAP* pipeline uses a running quartile filter window to determine the continuum. For this reason it is provably variant with the S/N ratio (biased at low S/N): at low S/N ratios the inferred continuum will be systematically shifted with respect to the same star observed in high S/N. Although our reference set only includes high S/N spectra, our results would suffer if we employed the *ASPCAP* normalization procedure for low S/N spectra in the *test set*. Thus, we opted to normalize and stack the fluxes from individual visits provided in the `apStar` files. Our normalization procedure is described below.

Individual spectra in the `apStar` files contain associated uncertainty arrays for each observation. However the uncertainty arrays do not encapsulate all knowledge about any technical, observational, or reduction issues. Instead, every pixel in a single observation also contains a bitmask flag that documents potential issues. Although pixels are flagged, these issues are not reflected in the uncertainty arrays of individual observations. For this reason we chose to construct an (adjusted) inverse variance array that encapsulates the implied additional uncertainty in flagged pixels. Specifically for pixel j, if it is flagged, we set the error array such that:

$$\sigma^2_{j,adjusted} = \sigma^2_j + \Delta^2_j \tag{6}$$

where

$$\Delta_j = \max\left(C_0|y_j - \widetilde{y}|, C_1\widetilde{y}N_{\rm flagged}\right). \quad (7)$$

Here $N_{\rm flagged}$ is the number of flagged pixels in the spectrum, and (what we call the conservatism) constants C_0 and C_1 have been chosen as $\{2.0, 0.1\}$, which produces sensible inverse variance values for flagged pixels. We chose to ignore bitmask values associated with persistence (specifically values 9, 10, and 11) as frequently every pixel in the blue CCD would be flagged with these values, thereby producing unrealistically large uncertainties at every pixel in the blue CCD. Consider the two limits of (7): pixel A is flagged and has a flux value that is very far (say 5σ) away from the mean flux. Pixel B is also flagged but has flux values very similar to the mean spectrum – a visual inspection suggests it is not anomalous. For flagged pixels of the A kind, the added variance is proportional to how discrepant the flux is to the mean spectrum flux. On the other hand, there may be many B-like pixels that represent a subtle systematic problem. These pixels are accounted for by considering the number of flagged pixels $N_{\rm flagged}$. In both cases, our conservative approach accounts for underestimated variance in flagged pixels due to local (A-type) and global (B-type) problems.

After updating the inverse variance arrays to account for flagged pixels, we pseduo-continuum-normalized the individual observations. The spectra were normalized in three different regions, corresponding to each CCD, between wavelength regions 15090–15822 Å, 15823–16451 Å, and 16452–16971 Å. We distilled a list of continuum pixels (following the initial identification in Ness et al. 2015a), and for each region in all observations we fit the continuum-pixel fluxes as a sum of sine and cosines with $L = 1400$ and $W = 3$, solving for amplitudes $\boldsymbol{A}$ in:

$$y_j = \sum_{w=0}^{W} A_{2w} \sin\left(\frac{2w\pi\lambda_j}{L}\right) + A_{(2w+1)} \cos\left(\frac{2w\pi\lambda_j}{L}\right). \quad (8)$$

Using a sum of sines and cosines for pseudo-continuum normalization has a number of advantages over alternative approaches. Firstly, a design matrix can be constructed *a priori*, at which point the normalization procedure becomes a linear matrix operation on fluxes and is therefore cheap. More practically, the use of sine and cosine functions implies that edge behaviour of the continuum function will be bounded (more bounded than a polynomial expression), and are therefore less susceptible to issues with overfitting where the edge of the continuum function demonstrates large wiggles.

The approach we have described is *pseudo*-continuum normalization. While our choice of continuum pixels is well-informed (Ness et al. 2015a), our procedure does not require

that the chosen pixels to are indeed "true" continuum pixels. This is advantageous. The continuum pixels could be randomly selected; our approach would be equally effective as long as the same pixels were always used for normalization. All spectra would be "normalized" in the same way, and the residual flux behaviour from a "true"-normalized spectrum would be captured by the spectral model. We have attempted to select well-informed continuum pixels only in order to maximize model interpretability: by using reasonable continuum pixels we can be sure that the spectral model derivatives can be interpreted as being astrophysically motivated, rather than capturing residual continuum effects.

For all of the aforementioned reasons (invariance with respect to S/N, bounded functions, linear operations, and repeatability between spectra), we emphasize that if traditional continuum methods are employed, *The Cannon* will (likely) give *very bad results*. Adopting a linear continuum normalization procedure is paramount. After normalizing all individual spectra in the `apStar` files, we re-stacked the spectra using the (adjusted) inverse variance of each pixel as weights. Therefore for all stars observed by *APOGEE*, we have normalized individual and combined spectra, allowing for a self-consistent examination of label determination at low S/N (see Section 4.3).

4. Experiments

4.1. Hyper-parameter selection and validation

In the simplest case we have two hyper-parameters that need to be determined: Λ and f. The strength of the regularization (at all pixels) is set by Λ, and the scale factor f controls the scaling on individual labels. In practice large Λ values encourage zero coefficients and thus enforce sparsity. High scale factor f values act to penalize cross-terms more than linear labels, and produce sparse models as a consequence. Because these hyper-parameters act in conjunction with each other, we need to explore many combinations of Λ and f.

Here we use the reference set described in the previous Section to perform a grid search with different values of Λ and f. Specifically we vary Λ in logarithmic steps from 10^0 to 10^5, and increase the scale factor f from 0.5 to 50. For each combination of Λ and f we trained the full 17-label model using the training set (the $\sim$90% random subset of the reference set) and measured labels for all spectra in the validation set: both the combined and individual observations. Given the model complexity, for this experiment we fixed $s^2 = 0$ to leverage convexity and to ensure convergence at the global minimum. This also ensures a fair comparison between different hyper-parameters, because if s^2 is free it will inevitably rise at very high Λ values.

We must chose heuristic(s) to select an appropriate regularization and scale factor. Many metrics are available to us: the predictive power in spectral fluxes (e.g., the χ^2 or test scalar for the validation set), the recovered precision in labels at low S/N, or something that balances goodness of fit and sparsity.

We define the model sparsity as the percent of zero-value spectral derivatives $\boldsymbol{\theta}$. The sparsity could be calculated across the linear coefficients, only the cross-term coefficients, or some combination thereof. Note that we never include the baseline spectrum coefficients $\boldsymbol{\theta}_0$ when calculating sparsity metrics because this parameter is not regularized (see Equation 5). In Figure 2 we show three different sparsity metrics for many permutations of hyper-parameters Λ and f. Specifically we show the sparsity of the linear model coefficients $\boldsymbol{\theta}_{1...17}$, the cross-term coefficients $\boldsymbol{\theta}_{18...170}$, and the combination of linear and cross-term coefficients. It is clear that the total model sparsity does not change significantly (regardless of f) until $\Lambda \gtrsim 10^3$. In this regime the linear and second-order coefficient sparsity metrics exhibit very different responses. The cross-term sparsity increases faster than the linear terms, with a clear (and expected) dependence on the scale factor f. For any $\Lambda \gtrsim 10^3$, a scale factor $f \approx 20$ produces the sparsest model.

While sparser models are preferred, our ultimate goal is to have an interpretable model that predicts spectral fluxes and returns precise stellar labels. In Figure 3 we show the median absolute deviation in abundance labels, measured between individual and combined spectra for stars in the validation set. This is an internally-consistent check for label recovery at low S/N: we will validate our high- (and low-) S/N label determination against *ASPCAP* values in Section 4.2. It is clear that a combination of decreasing f with increasing Λ recovers labels for validation set stars with high precision. At $\Lambda \approx 10^3$, varying f between 0.5 and 50.0 results in a marginal change in the validation set precision (0.04 dex to 0.06 dex). Figure 3 suggests that at $\Lambda = 10^3$, $f = 0.5$ has comparable performance in label recovery as a model with the same Λ and $f = 50.0$. While this behaviour was observed in the recovery of some abundance labels, it was not seen in all. The behaviour differed for each label. However, lower scale factors were favoured by all labels.

As a final heuristic to guide our choice of Λ and f, we examined the performance in predicting spectral fluxes for all validation set spectra. Here we predicted stellar fluxes for all validation set stars (using the *ASPCAP* labels) and calculated the total χ^2 difference. The results are shown in Figure 4 compared to the χ^2 for the least-regularized case ($\Lambda = 10^0$). At increasing regularization strength the models demonstrate better predictive power in spectral fluxes, reflected by a lower total χ^2 value. As expected, the χ^2 minimum is dependent on the scale factor. While regularized models with lower χ^2 values clearly predict spectral fluxes more accurately (a comparison with *ASPCAP* predicted fluxes is shown in Figure 5), this

heuristic only gives a weak limit on our choice of hyper-parameters: *any* regularized model with a lower χ^2 value than seen in the unregularized case is objectively a better model because it predicts stellar spectra more accurately.

Our grid search has revealed that a regularization factor of at least $\Lambda = 10^3$ is required to produce a sufficiently sparse model. Generally we find that low scale factors yield marginally better behaviour in recovering abundance labels at low S/N. At $\Lambda = 10^3$, the precision in abundance labels would suggest that a scale factor anywhere between $f = 0.5 - 5.0$ is reasonable. However the total χ^2 for validation set fluxes shows $f = 2.0$ to be a far better model (in terms of predictive power in spectral fluxes) than $f = 0.5$ or $f = 5.0$. Thus, on the basis of these metrics, we adopt $\Lambda = 10^3$ and $f = 2.0$ as our *regularized model* hyper-parameters for the remainder of this work.

4.2. Validating the regularized model

We re-trained the 17-label model (without requiring $s^2 = 0$) using our adopted hyper-parameters. In this and subsequent sections we validate our model through various experiments. Our first test is to examine how well we can recover labels for the validation set stars. Recall that these stars are a random 10% subset of the reference set and were not used for training. Comparisons between the *ASPCAP* validation set labels and our own are shown in Figure 6, where we find that the model shows good agreement with the *ASPCAP* labels. No strong biases are seen, and the standard deviation of the residuals is low for most abundance labels. The largest disagreement is seen in [V/H] and [Na/H], in different regimes of label space. While [Na/H] starts to become discrepant at [Na/H] < -1 (perhaps because we are reaching the noise floor for this label), [V/H] becomes increasingly discrepant near [V/H] ≈ 0.

The [V/H] abundance labels are problematic in the *ASPCAP* pipeline because the signature is subtle even at high S/N ratios: there are very few spectral lines, most of which are blended with stronger lines. Indeed, we noted very obvious issues in the [V/H] labels in the reference set which we tried to account for (e.g., Section 3). Although our [V/H] label precision is worse than other abundance labels, it is limited by the training set.

We report formal errors in all labels provided by *The Cannon*. At test time the Jacobian matrix of partial derivatives is estimated at the optimized solution, which we multiply by the residual variance to obtain a formal covariance matrix. Errors in each label are taken as the square-root of the diagonal entries along the covariance matrix. We stress that the errors listed in Table 2 are formal errors only: they do not encapsulate all aspects of the error

budget that are known (or unknown) to us. An uncertainty floor in each label is estimated in the following Section, which ought to be added in quadrature with the formal errors.

4.3. Label recovery as a function of signal-to-noise

Here we seek to understand the capability of our regularized model to recover labels as a function of S/N. The validation set includes stacked spectra for each star as well as the individual normalized observations. We measured labels for every spectrum. This permits us to understand how well we can recover labels (to some precision) for a given S/N. Crucially, this experiment provides a very model-independent metric for performance. At low S/N ratios the label errors will be dominated by random uncertainties. In the high S/N regime, systematic uncertainties are more relevant. Knowing exactly where this "turnover point" occurs from systematic to random uncertainties is a useful metric to compare analysis pipelines.

The results of our label recovery experiment from *individual visit* spectra of validation set stars are shown in Figure 7. At $S/N \gtrsim 50$ the discrepancies in all labels are either flat or taper towards zero, indicating that we are dominated by systematics. This is an important transition point, because *stacked* spectra for all bonafide giant stars (Section 5) have $S/N \geq 50$. In this regime we recover labels from within 0.01 dex ([Fe/H]) to 0.16 dex ([V/H]) for all abundance labels. The worst performance is again found for [V/H], by a factor of two: the second-worst label at $S/N \geq 50$ is [Na/H] with 0.08 dex. These values are listed for all labels in Table 1, where we advocate their use as a representative error floor that should be added in quadrature with the formal errors. We stress that this is not a thorough accounting of the error budget.

It is clear that the regularized model recovers labels with good precision even in the presence of substantial noise. In our experience the recovered precision presented here is higher than most physics-driven analysis pipelines. The reason for this is not completely understood (by us), because both approaches (optimizing a data-driven or physics-driven model) rely on least-squares fitting, which will usually follow the $1/\sqrt{S/N}$ behaviour we see in Figure 7 unless there are strong non-linearities in the model. Unfortunately we cannot show the same comparison for *ASPCAP* labels as a function of S/N in Figure 7 as *ASPCAP* labels are not publicly available for individual visit spectra.

5. Results

Our experiments have demonstrated that a data-driven model for stellar spectra can be reliably extended to high dimensionality in label space. We have further shown that the regularization hyper-parameters can be simplified to just two hyper-parameters that can be set heuristically. This yields a sparse, interpretable model that recovers labels with high precision at low S/N. While this is reassuring, it is less relevant for our results as all stacked *APOGEE* spectra for bonafide giants (see below) have $S/N \gtrsim 50$, well into the regime where we are advantageously dominated by systematic uncertainties.

We have used our regularized model to measure (test) labels of 150,677 *APOGEE* spectra, all normalized and stacked using the method in Section 3. In addition to the model being effective, the test step is very fast: our pure-`Python` implementation returned *17 labels for all 150,677 spectra in just 28 minutes* of wall-time on a small research cluster in Cambridge. These were free and otherwise unused resources; no dedicated computing assets were required. This pace is also projected to increase, as the test step did not include analytic derivatives $\delta\boldsymbol{\theta}/\delta y_j$, which are now implemented in our open-source code for any polynomial vectorizer. The test-step optimization is not convex because the vectorizer contains quadratic label terms. For this reason we ran the optimization from nine different initialization points, chosen to sparsely cover the range of $T_{\rm eff}$, $\log g$, and abundance labels in the training set. Of the nine optimizations, we adopted the end result with the lowest χ^2 value.

The training set only includes giant stars, but the *APOGEE* DR12 includes giants and dwarfs. Therefore we exclude results with $\chi^2_r > 3$, stars without a $T_{\rm eff}$ label from *ASPCAP*, or stars with *ASPCAP* labels outside of the bounds $5500 > T_{\rm eff} > 4000$ or $\log g < 3.9$. The distilled sample contains 87,563 giant stars, where we report $T_{\rm eff}$, $\log g$, and 15 abundance labels. The distribution of χ^2 values for all 150,677 combined spectra are show in Figure 8. The labels in the distilled sample follow expectations from stellar astrophysics, and include stars that are marginally outside the training set. For example, the *ASPCAP* labels include a strict cut in $T_{\rm eff}$ at 3600 K, but we reliably recover labels beyond this boundary. Figure 9 presents a few different label projections for the distilled sample, indicative of the boundaries and distribution of our labels. The full complement of 17-labels for 87,563 stars is provided in Table 2, in addition to quality metrics.

Our abundance labels are consistent with detailed studies of galactic chemical evolution. In Figure 10 we show our abundance labels (with respect to iron) of α-capture (e.g., O, Mg, Ca, Si), odd-Z (Al, Na), light (C, N), and Fe-peak (Ti, Mn, Ni) elements for all bonafide giant stars in *APOGEE* DR12. These abundances trace nucleosynthetic pathways from different sources (e.g., supernovae, AGB stars) and reflect the environmental conditions at the time of their formation. In Section 6 we will briefly discuss the implications and interpretations

of these abundance labels.

We trained *The Cannon* using high-fidelity labels from *ASPCAP*. As expected, our labels agree excellently with *ASPCAP* for stars with high S/N ratios. This is not the case for stars with S/N ratios below 200. Here we present abundance label projections for some astrophysically interesting subsets of the *APOGEE* sample, and show comparisons between *The Cannon* and *ASPCAP*. Our first comparison is shown in Figure 11, where we have selected a sequence of stars with high [α/Fe] abundance ratios and shown three well-behaved (precise) abundance labels that trace different nucleosynthetic pathways. Points are coloured by their S/N ratio, and the black lines indicate density contours.

The differences between *The Cannon* and *ASPCAP* projections in Figure 11 are striking. Labels from *ASPCAP* have a much larger range than those from *The Cannon*, however nearly all stars that fill this difference in label range are those with low S/N ratios. This implies that either those stars with lower S/N are sampling a different (presumably more distant) part of the Milky Way, or the precision in *ASPCAP* labels degrades much worse at low S/N ratios than *The Cannon*. Our validation tests have demonstrated that our label precision remains approximately constant (systematic-dominated) at $S/N \gtrsim 50$ (i.e., for all combined spectra of bonafide giants in *APOGEE* DR12). For these reasons we argue that the difference in abundance projections between *The Cannon* and *ASPCAP* for low S/N stars in the [α/Fe] set are likely because *ASPCAP* yields imprecise results for spectra with S/N below $\sim$200. Although the imprecise *ASPCAP* results can presumably be discarded with some quality cuts, a substantial fraction of the *APOGEE* sample would have to be removed.

Globular clusters are excellent laboratories for us to validate our abundance precision and label interpretability. Indeed, it is our firm belief that there is sufficient new information *just* in our globular cluster abundance labels that would warrant its own publication. We encourage others to pursue this; here we simply comment on how our results compare to *ASPCAP*, and in the next Section we will discuss how our results compare to existing studies on globular clusters, and expectations from stellar evolution. Here we discuss three globular clusters that have a considerable number of bonafide members in DR12: M 15, M 53, and M 92.

In each cluster we attributed stars to be bonafide members based on the *ASPCAP* radial velocity and *The Cannon* [Fe/H] label. The top panel of Figures 12-14 shows our membership selection with respect to all other *APOGEE* giants in the same field (i.e., those that share the same '`FIELD`' entry). Although our membership criteria may suffer from mild contamination of field stars, this does not significantly impact our comparisons to *ASPCAP* labels or any discussion in the following Section. The remaining panels of Figures 12-14 show individual abundance projections using *ASPCAP* labels (left) and those from *The Cannon* (right). In

any projection, the abundance labels from *The Cannon* show a smaller spread than *ASPCAP* for the same stars. In most cases the mean cluster labels between *ASPCAP* and *The Cannon* are similar, with correlations in the same direction. The only difference is that smaller label distributions are found by *The Cannon*. The most striking comparison is seen for M 15 (Figure 12).

6. Discussion

We have demonstrated that a data-driven model for stellar spectra can be reliably extended to high dimensionality in label space. We return high precision labels for *APOGEE* spectra across the giant branch, including those with low S/N ratios. For these noisy data, the abundance labels returned by our model for low S/N ratios show similar behavour to what is seen in the high S/N (> 250) sample. However our abundance labels from noisy data are significantly different from *ASPCAP*. In this noisy regime there are good reasons to believe *The Cannon* abundance labels are more reliable. This has a prominent impact on abundance labels for globular cluster stars, as our results suggest much narrower intrinsic abundance spreads ($\sim$3-4 times smaller).

There are (at least) two possibilities to explain the discrepancies in cluster abundances between *ASPCAP* and *The Cannon*. The first is that there is a large underlying intrinsic abundance spread in these elements for each cluster, and – for whatever reason – *The Cannon* is returning systematically inaccurate abundances. The alternative is that the intrinsic cluster abundance spreads *are* smaller than what *ASPCAP* labels suggest, and the larger cluster spreads in *ASPCAP* abundances are simply due to large uncertainties in individual abundances. On the basis of our validation tests, we argue the latter to be a more plausible explanation for these abundance discrepancies.

The overall metallicities ([Fe/H] label) we find for the globular clusters examined are in satisfactory agreement with the literature. In general our [Fe/H] label is more metal-rich than what is quoted in existing cluster compilations (Harris 1996). The difference seems to be an imprint from the *ASPCAP* labels at low-metallicity, as the *ASPCAP* labels for stars we assign as cluster members are also more metal-rich than literature sources. This deviation may itself be linear in the *ASPCAP* [Fe/H] scale, as the agreement is very good for metal-rich clusters in our sample. For example, for M 3 we find a cluster mean and standard deviation of [Fe/H] $= -1.38 \pm 0.09$, just 0.12 dex more metal-rich than the literature mean (Harris 1996). Similarly we find M 13 to have [Fe/H] $= -1.45 \pm 0.09$ whereas the same source quotes [Fe/H] $= -1.53$. However at the metal-poor end the disparity becomes severe, where we find a 0.5 dex offset with the literature for M 92 ([Fe/H] $= -1.81 \pm 0.04$).

While these mismatches are important to examine, our tests suggest they are a reflection of two compounding issues: a relative paucity of metal-poor stars in the training set, and a suggestion of a linear-in-[Fe/H] systematic trend in *ASPCAP* [Fe/H] labels. A more thorough comparison of metal-poor stars in *ASPCAP* may help resolve these issues. Nevertheless, while *The Cannon* abundance scale may be slightly inaccurate (offset from the literature consensus) for metal-poor stars, this issue does not impact our precision in abundance labels that we report for these clusters.

The anti-correlation we see between [C/Fe] and [N/Fe] abundances in M 15 is expected from the CN-cycling of stellar material. The labels (and correlation strength) are in excellent agreement with high-resolution studies on this cluster (Cohen et al. 2005). Indeed, while we find an intrinsic spread in these abundance labels, we find that the spread is much smaller than considered previously. We also corroborate works that find no evidence of bi-modality in C and N abundance labels, a signature commonly found in metal-rich globular clusters. While the *ASPCAP* labels show a substantial spread in total C+N+O abundances for M 15, our results imply that CNO-cycling has maintained an approximately constant sum abundance of these elements, consistent with recent works (Mészáros et al. 2015).

The reference set described in Section 3 was not purposefully constructed to include stars in *any* globular cluster. We selected stars that met strict quality criteria, and discarded stars with questionable label fidelity. Given the abundance (anti-)correlations we confirm from other high-resolution studies, our construction of the reference set has a number of implications. First, *the anti-correlations seen in globular clusters here and elsewhere are not dominant in our training set*. Indeed, the lack of globular cluster (metal-poor) stars in our training set partially explains the [Fe/H] discrepancy we see with respect to the literature (i.e., very few metal-poor stars in the training set). However as a consequence the regularized model is successfully measuring sets of labels (e.g., chemical fingerprints) that are *very different* to what the model was actually trained on. Even though these chemical signatures are distinct from the training set, an experimental astrophysicist expects to see them based on previous works. For this reason it is very pleasing to see these signatures because (amongst other things) it reflects the interpretable nature of our model: the spectral derivatives *do* have physical meaning, allowing for the model to recover label patterns of high astrophysical interest that are substantially different from the training set.

Our measured labels have good astrophysical interpretations that are consistent with expectations from stellar evolution. The photospheric abundances of [C/H], [N/H] and [O/Fe] are expected to change as a giant star experiences CNO cycling. This pattern is present in our data, and to a lesser extent it is visible in the *ASPCAP* labels. Whilst these abundances vary, the sum of these elemental abundances is expected to remain roughly

constant for a giant star, regardless of its evolutionary state. This is not seen in the *ASPCAP* labels in any globular cluster described here. A huge spread in total [(C,N,O)/3Fe] abundance labels is seen, which is perhaps reflective of the sum of large uncertainties in each *ASPCAP* C, N, and O abundance label. In contrast, *The Cannon* labels remain roughly constant for each globular cluster.

In addition to returning precise abundance labels, the regularized model we have presented is interpretable. Machine learning techniques can be extremely useful, but they are frequently limited in their application because they usually do not have internal components that are interpretable. This is not the case for *The Cannon*: the spectral coefficients $\boldsymbol{\theta}_0$ represent the baseline spectrum, and the first-order coefficients ($\boldsymbol{\theta}_{1-17}$ in this model) represent the spectral derivatives with respect to each label. They indicate where specific labels (or elements) contribute to the spectrum. In Figure 15 we demonstrate an example of this, where the first-order derivatives of three sparsely-acting labels [Al/H], [S/H] and [K/H] are shown for two spectral regions. The locations of known strong atomic lines are indicated by vertical markers (Smith et al. 2013). Although no information about atomic line lists enter into our spectral model, the spectral derivatives for [Al/H] and [K/H] show strong contributions at wavelengths of known atomic lines.

The two unmarked spectral lines in Figure 15 correspond to atomic lines that are seen in *APOGEE* spectra, but for which there is currently no atomic data available. The element that is responsible for these spectral lines is unknown (Shetrone et al. 2015). Figure 16 shows a zoom-in around these two regions, where we show *all* first-order spectral derivatives and color [S/H] and [Al/H] as per Figure 15. While there is considerable blending of spectral derivatives near the line at 15235 Å, [S/H] shows the dominant first-order spectral coefficient at all pixels that include this (previously unknown) atomic line. For the previously unknown atomic line at 16755 Å, the spectral derivatives in our regularized model strongly suggest that this is an Al transition. At 16755 Å the [Al/H] spectral derivative is nearly the strongest seen anywhere in the *APOGEE* spectra, matched only by two nearby (known) Al lines, and there is no significant blending by spectral derivatives of other abundance labels.

We made no effort to cross-match our model with the list of "unknown" spectral lines identified by the *APOGEE* group. This example presented itself serendipitously whilst examining the relative strengths of different first-order derivatives. Indeed, the list of "unknown spectral lines" was unknown to the lead-author when Figure 15 was first produced. While it appears likely that the two unknown lines shown here arise from S and Al, respectively, we stress that the spectral derivatives should only be used as a guide. The abundances of different elements are strongly covariant in nature (see below). For this reason the spectral derivatives are merely indicative as to which element label (of the set that we have available)

correlates most with the fluxes at the given pixels. We have not performed a detailed examination to understand *how interpretable* this *Cannon* model is, and what physical insights it can provide. Convincing examples of this covariant behaviour (and methods to address it) will be presented in a companion paper (Ness et al., 2016, in preparation).

Indeed, there is considerable room for improvement over the work we have presented here. For example, at least one of our model assumptions is provably false. *APOGEE* spectra have different resolutions. The line spread function is both wavelength- and fibre-dependent. Given this information, our current implementation of *The Cannon* is sub-optimal: we assume nothing about differing line spread functions between stars in the reference or test set. Similarly, we make no effort to accommodate fast-rotating stars, where the effect on the spectrum is approximately represented by convolution with a Gaussian kernel (in the same way a lower resolution would be). Tests performed by M. Ness have revealed that this information is present in the data: if the fibre number is included as a label in ℓ, then the fibre number can be accurately inferred from stars in the test set! This is possible because the *APOGEE* resolution varies smoothly (to some degree) with fibre number. For these reasons it is clear that simultaneously solving for additional rotation at test time would be an useful extension of this work.

Notwithstanding these issues, *The Cannon* clearly has its place amongst the current methods employed to analyse stellar spectra. Precise labels are returned even in the presence of noisy data, and the model internals are sufficiently interpretable that they can inform physics-based models. Admittedly, our comparisons to the globular cluster literature suggests that our [Fe/H] (and other?) labels may be inaccurate in the metal-poor regime, a signature arising from incompleteness and inaccuracy of our training set labels. However we have demonstrated that this limitation does not hinder any potential inferences using individual chemical abundances: for these (and many other purposes) we care more about abundance precision, not accuracy. With increased precision, we have shown that globular clusters exhibit significantly narrower intrinsic spreads in abundances (and their anti-correlations) than previously considered. Crucially, because the detailed chemical abundances reflect the star forming conditions and subsequent evolution of the cluster, this work is a necessary and fundamental step forward in understanding the evolution of stars, clusters, and ultimately, the Milky Way.

All of the code for this project is available with documentation at `http://thecannon.io/`. We encourage open community discussion by way of GitHub issues, and code contributions through pull requests.

The authors warmly thank Daniel Foreman-Mackey (University of Washington), Jason

Sanders (Cambridge), and Angus Williams (Cambridge) for valuable discussions.

This project was funded in part by the European Research Council under the European Union's Seventh Framework Programme (FP 7) ERC Grant Agreement 320360, the NSF (grants IIS-1124794, AST-1517237), NASA (grant NNX12AI50G), and the Moore-Sloan Data Science Environment at NYU. This research made use of: the NASA *Astrophysics Data System Bibliographic Services*, `TOPCAT` (Taylor 2005), GitHub, Travis CI, the `IPython` (Perez & Granger 2007), `numpy` (van der Walt et al. 2011), `matplotlib` (Hunter 2007), and `scipy` (Jones et al. 2001) packages, as well as Astropy, a community-developed core Python package for Astronomy (Astropy Collaboration et al. 2013).

This project made use of *SDSS-III* data. Funding for *SDSS-III* has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy Office of Science. The *SDSS-III* web site is http://www.sdss3.org/.

SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the *SDSS-III* Collaboration including the University of Arizona, the Brazilian Participation Group, Brookhaven National Laboratory, Carnegie Mellon University, University of Florida, the French Participation Group, the German Participation Group, Harvard University, the Instituto de Astrofisica de Canarias, the Michigan State/Notre Dame/JINA Participation Group, Johns Hopkins University, Lawrence Berkeley National Laboratory, Max Planck Institute for Astrophysics, Max Planck Institute for Extraterrestrial Physics, New Mexico State University, New York University, Ohio State University, Pennsylvania State University, University of Portsmouth, Princeton University, the Spanish Participation Group, University of Tokyo, University of Utah, Vanderbilt University, University of Virginia, University of Washington, and Yale University.

REFERENCES

Alam, S., Albareti, F. D., Allende Prieto, C., et al. 2015, ApJS, 219, 12

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, Astronomy & Astrophysics, 558, AA33

Cannon, A. J., & Pickering, E. C. 1912, Annals of Harvard College Observatory, 56, 115

Cohen, J. G., Briley, M. M., & Stetson, P. B. 2005, AJ, 130, 1177

De Silva, G. M., Freeman, K. C., Bland-Hawthorn, J., et al. 2015, MNRAS, 449, 2604

García Pérez, A. E., Allende Prieto, C., Holtzman, J. A., et al. 2015, arXiv:1510.07635

Gilmore, G., Randich, S., Asplund, M., et al. 2012, The Messenger, 147, 25

Harris, W. E. 1996, AJ, 112, 1487

Ho, A. Y. Q., Ness, M. K., Hogg, D. W., et al. 2016, arXiv:1602.00303

Hogg, D. W., Casey, A. R., Ness, M., Rix, H.-W., & Foreman-Mackey, D. 2016, arXiv:1601.05413

Holtzman, J. A., Shetrone, M., Johnson, J. A., et al. 2015, AJ, 150, 148

Hunter, J. D. 2007, Matplotlib: A 2D Graphics Environment, Computing in Science & Engineering, 9, 90-95

Jones, E., Oliphant E., Peterson P., et al. 2001, SciPy: Open Source Scientific Tools for Python, http://www.scipy.org/ (Online; accessed 2016-02-16)

Mészáros, S., Martell, S. L., Shetrone, M., et al. 2015, AJ, 149, 153

Ness, M., Hogg, D. W., Rix, H.-W., Ho, A. Y. Q., & Zasowski, G. 2015, ApJ, 808, 16

Ness, M., Hogg, D. W., Rix, H., et al. 2015, arXiv:1511.08204

Pérez, F., & Granger, B. E. 2007, IPython: A System for Interactive Scientific Computing, Computing in Science & Engineering, 9, 21-29

Shetrone, M., Bizyaev, D., Lawler, J. E., et al. 2015, ApJS, 221, 24

Smith, V. V., Cunha, K., Shetrone, M. D., et al. 2013, ApJ, 765, 16

Taylor, M. B. 2005, Astronomical Data Analysis Software and Systems XIV, 347, 29

Tibshirani, R. 1996, Journal of the Royal Statistical Society. Series B (Methodological), 58, 267-288

van der Walt, S., Colbert S. C., Varoquaux, G. 2011, The NumPy Array: A Structure for Efficient Numerical Computation, Computing in Science & Engineering, 13, 22-30

Zasowski, G., Johnson, J. A., Frinchaboy, P. M., et al. 2013, AJ, 146, 81

This preprint was prepared with the AAS LaTeX macros v5.2.

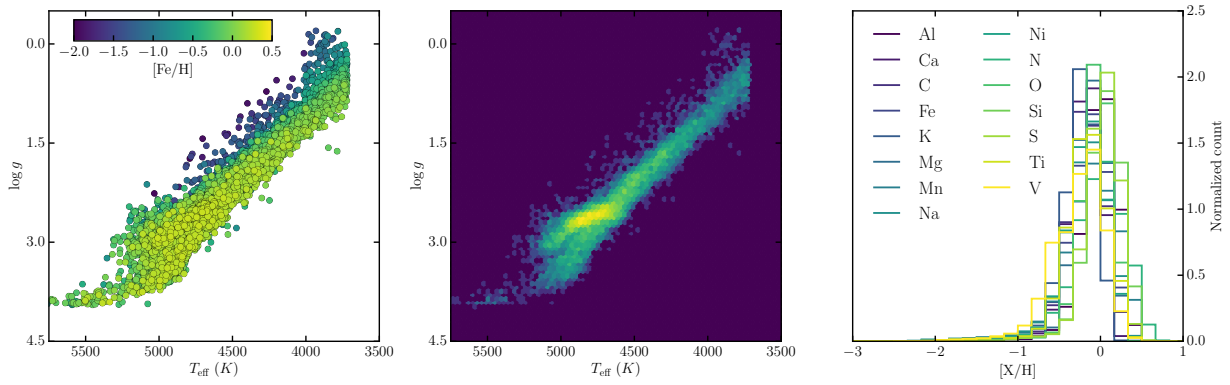

Fig. 1.— A temperature-surface gravity diagram for all 12,681 stars in the training set (left). A logarithmic density plot for the training set is shown in the centre panel. Distributions of abundance labels are shown in the right-hand panel.

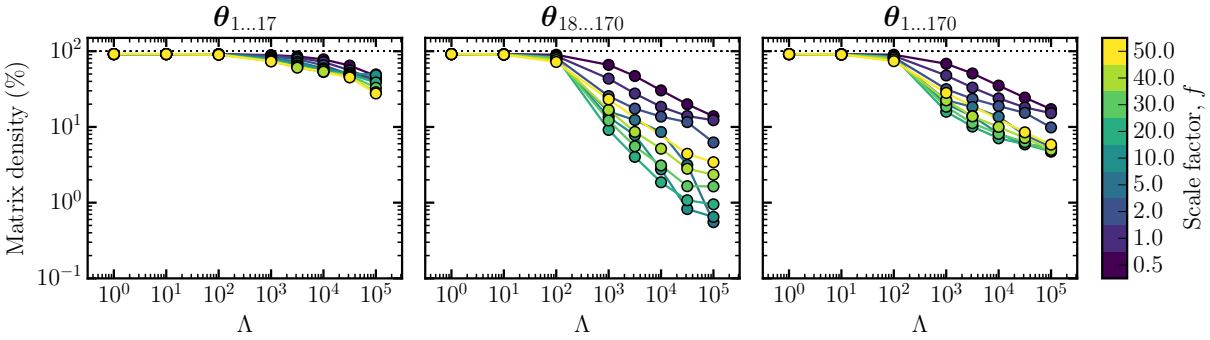

Fig. 2.— The fraction of non-zero coefficients (i.e., a measure of matrix density) for 17-label models with different hyper-parameters Λ and f. The density of the first-order coefficients are shown in the left panel, the second-order coefficients (e.g., $T_{\rm eff}^2$, $T_{\rm eff} \cdot \log g$, or [Al/H] $\cdot$ [Mn/H]) in the middle panel, and total density (excluding the baseline spectrum coefficient $\boldsymbol{\theta}_0$) in the right-hand panel.

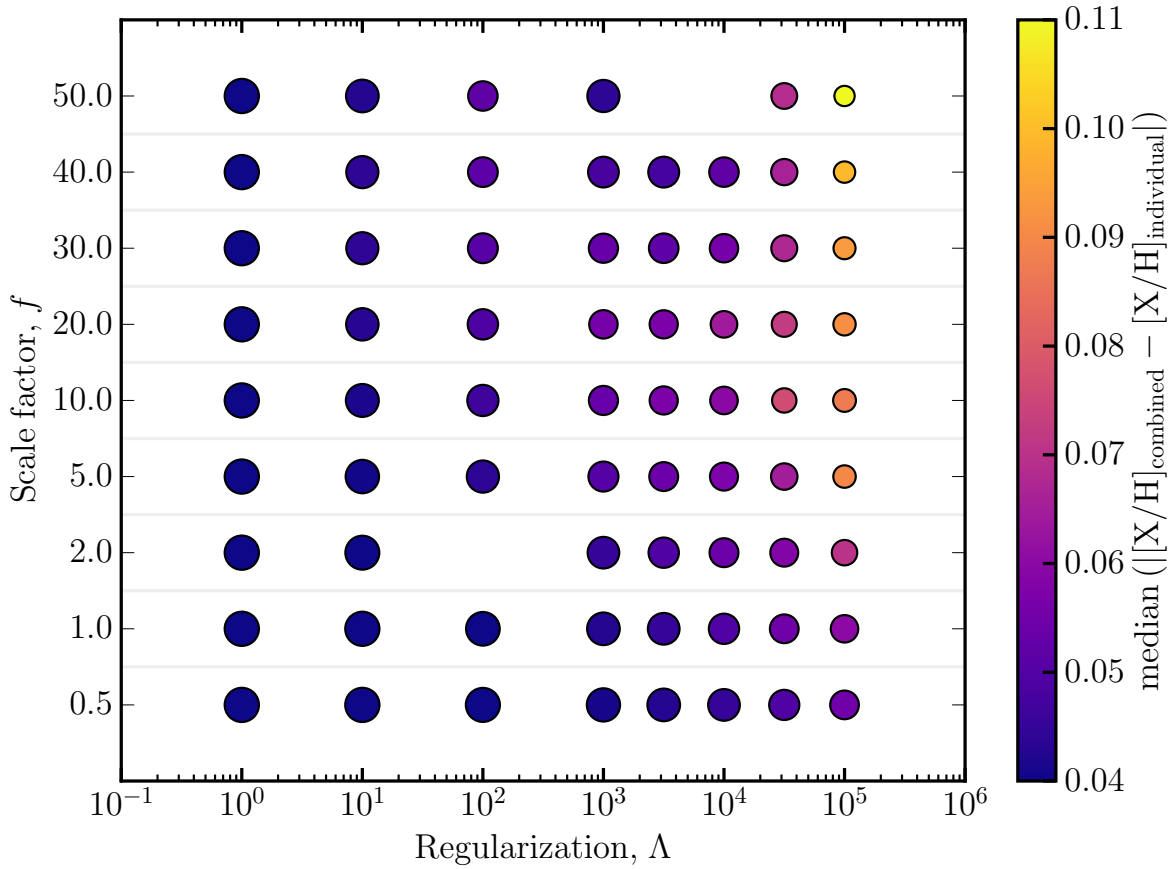

Fig. 3.— The median absolute abundance difference (over all abundance labels) between that inferred from individual observations in the validation set, and that inferred from the high S/N combined spectrum.

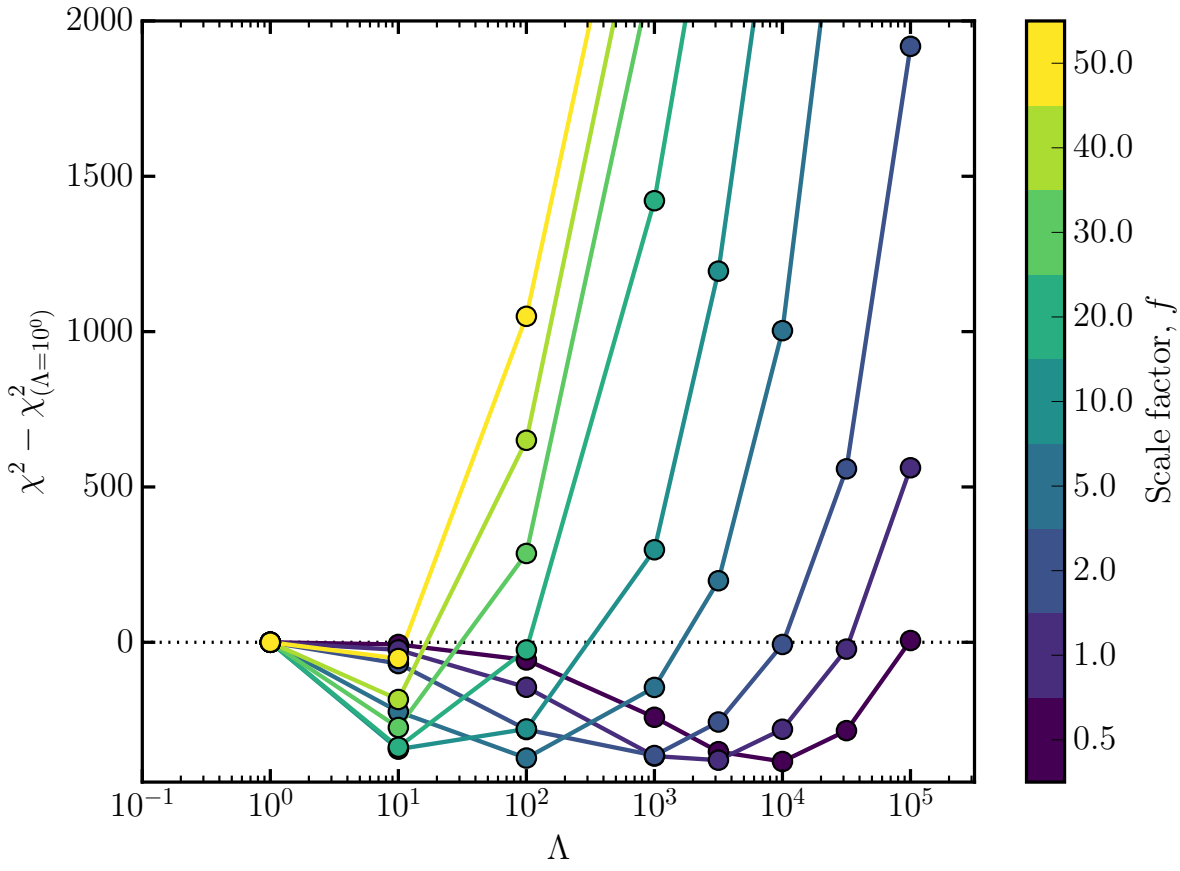

Fig. 4.— The total χ^2 (i.e., the *test scalar*) relative to the $\Lambda = 10^0$ model for all validation set stars for different combinations of the hyper-parameters Λ and f.

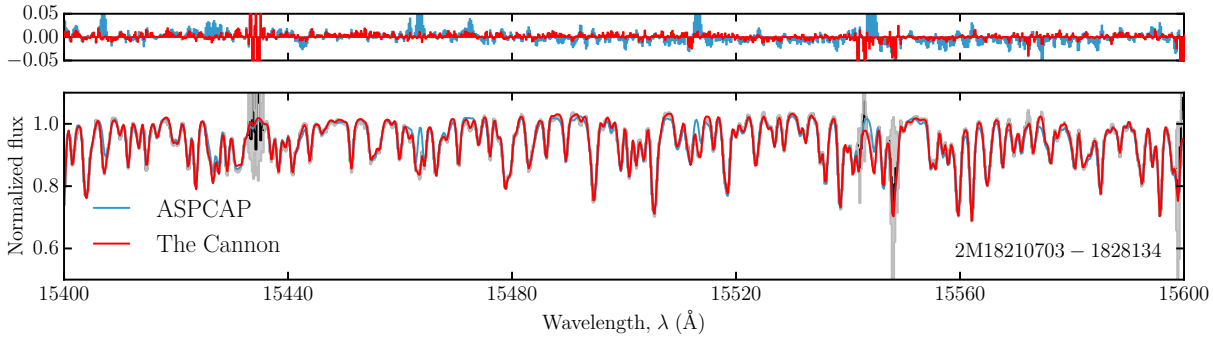

Fig. 5.— A portion of the (*Cannon*-)normalized spectrum for 2M18210703-1828134, a randomly selected star in the validation set. The best-fit model spectra for *ASPCAP* and *The Cannon* are shown. The residuals are plotted in the top panel.

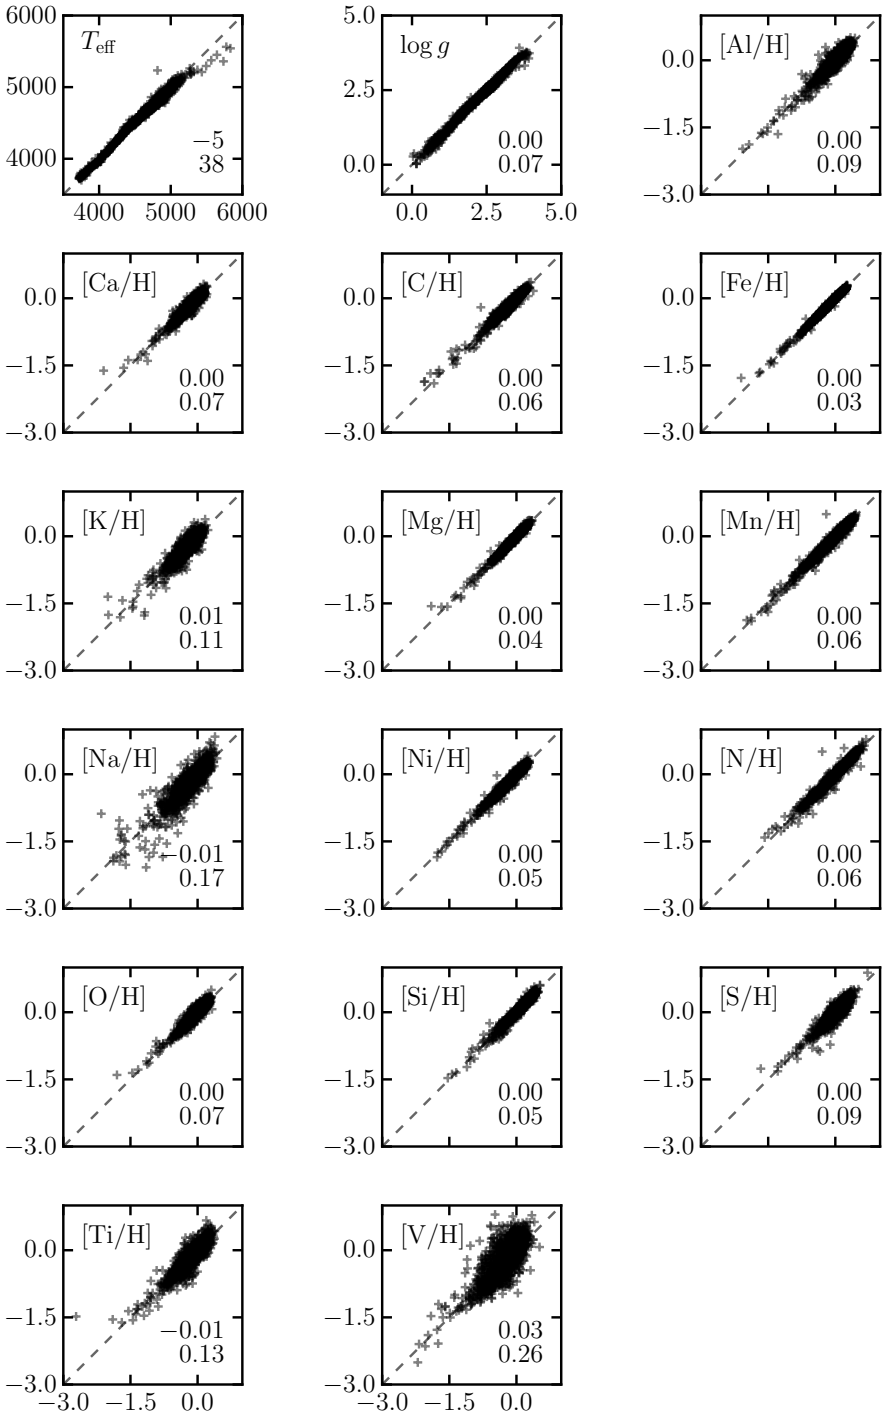

Fig. 6.— Comparisons between *ASPCAP* (x-axis) and *The Cannon* (y-axis) labels for the 1,460 high S/N stars in the validation set. For clarity purposes the label name is shown within each panel, as is the mean and standard deviation of the residuals. A dashed one-to-one line is shown.

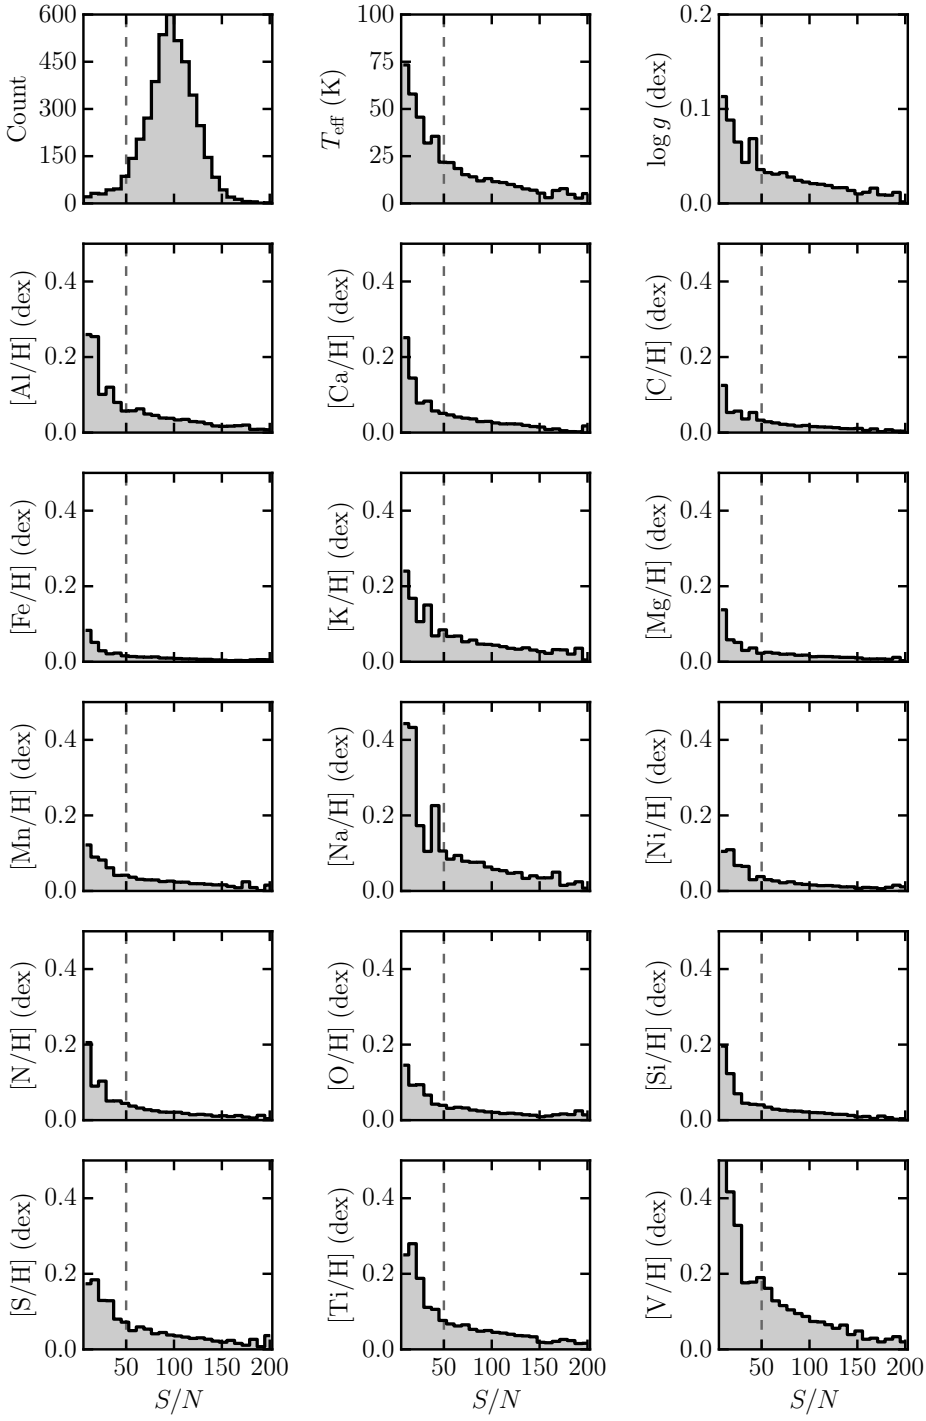

Fig. 7.— *Top left panel*: The distribution of S/N values for spectra of individual observations in the validation set. *All other panels*: The median absolute difference between the measured label for an individual observation, and the measured label from the high S/N combined spectrum for the same star. All panel have common bins, and all abundance labels have common axis limits. The dashed line indicates the minimum S/N for any *stacked APOGEE* spectrum in Table 2, implying the precision floor achievable for all *APOGEE* stars.

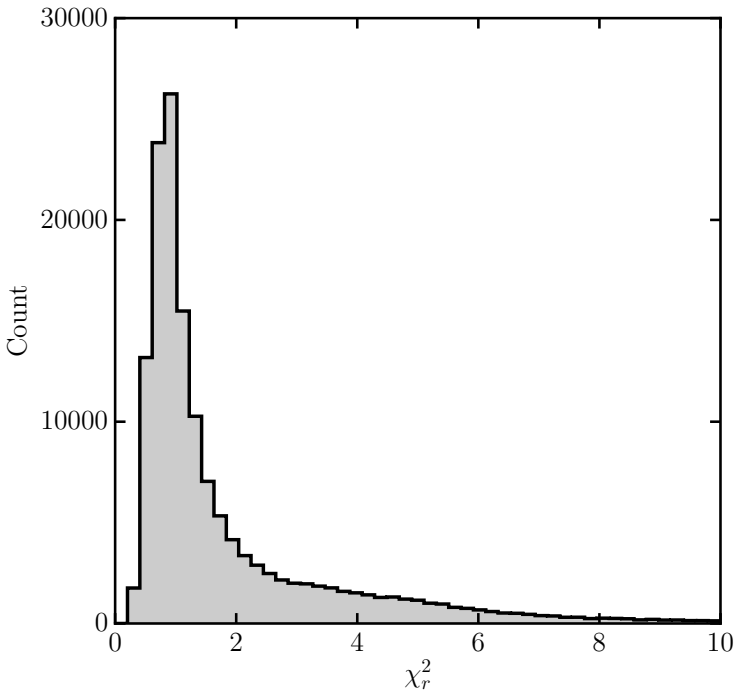

Fig. 8.— Reduced χ^2 values for all stars (stacked spectra) in *APOGEE* DR12.

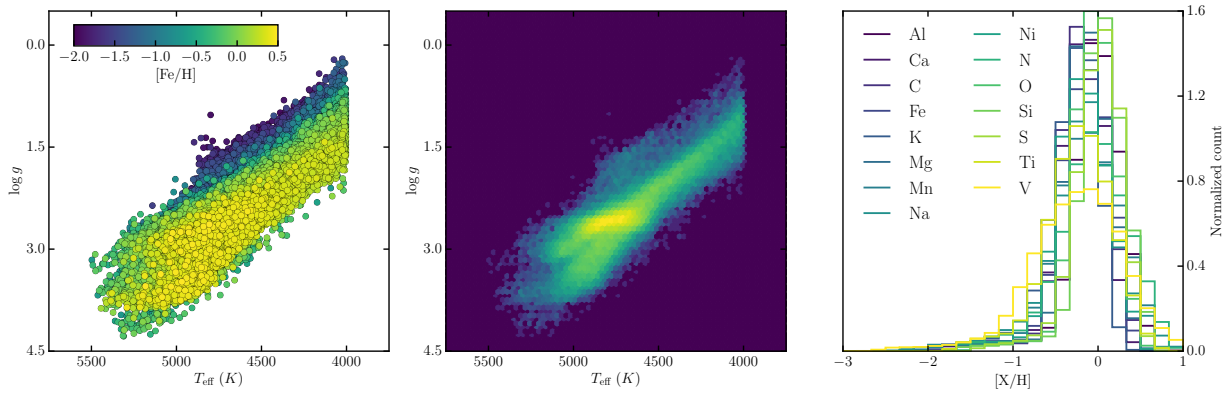

Fig. 9.— A temperature-surface gravity diagram for all 87,563 stars in the test set that pass our quality criteria (left). The centre panel shows a logarithmic density plot for the same stars. Abundance label distributions are shown in the right-hand panel.

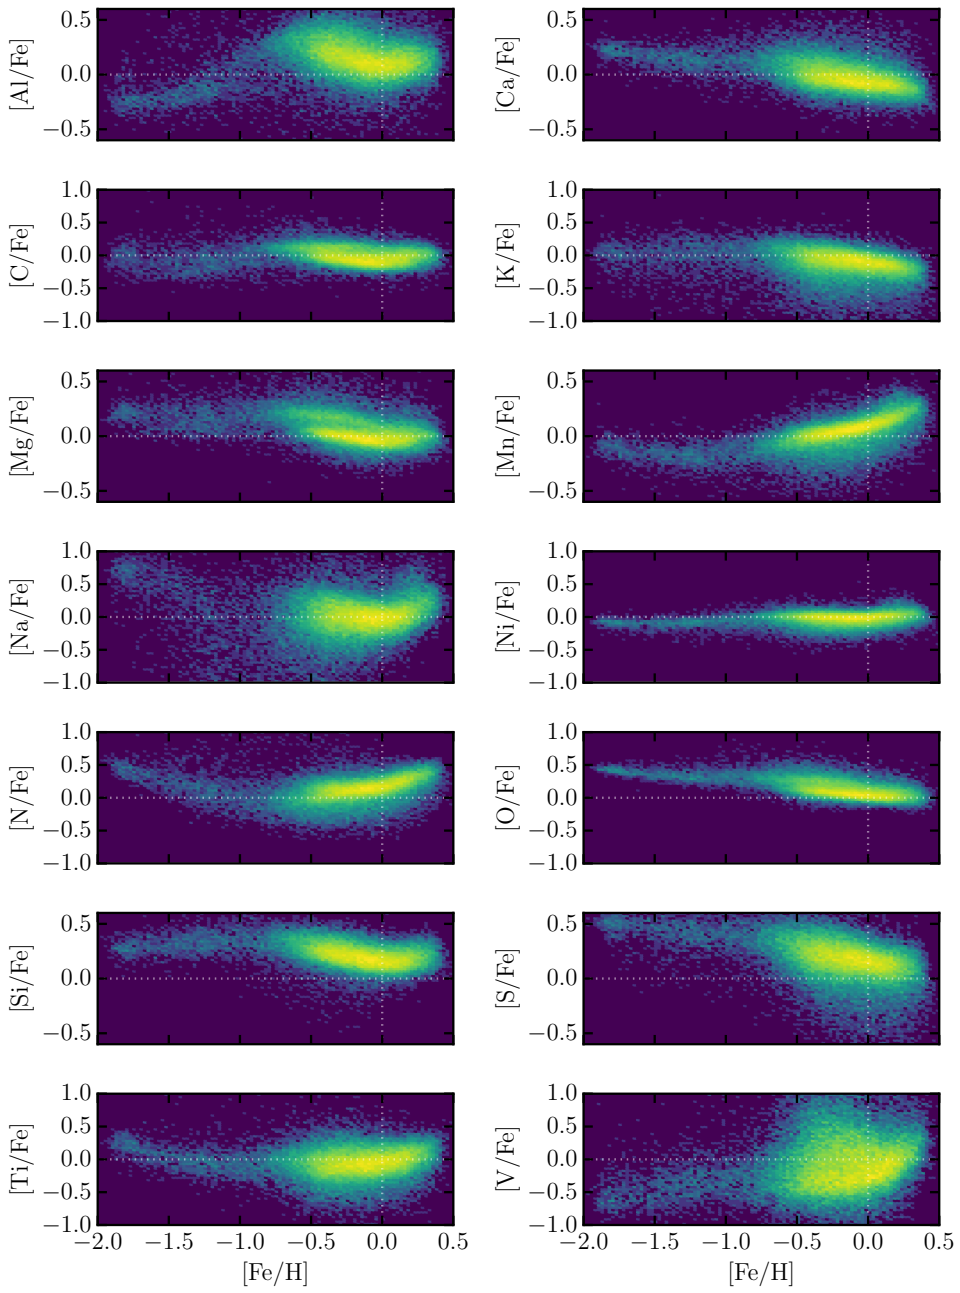

Fig. 10.— Logarithmic density for all abundance labels with respect to iron ([X/Fe]), revealing the Galactic enrichment with respect to Fe for all elements. Dotted lines indicate Solar abundances.

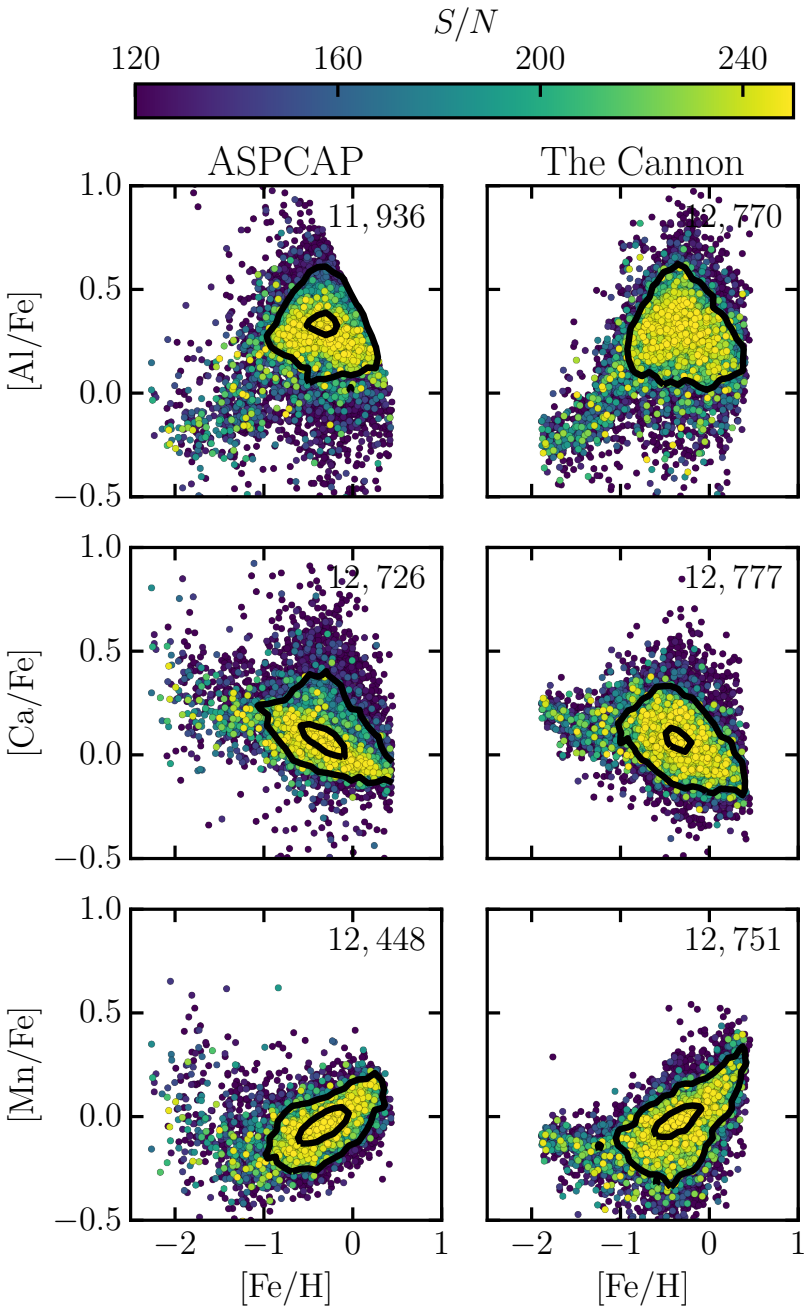

Fig. 11.— Abundance labels of three element groups (odd-Z, α, and Fe-peak) for stars in a high [α/Fe] abundance sequence. Labels from *ASPCAP* (left) and *The Cannon* (right) are shown for the same set of stars. Stars are colored by their S/N ratio to highlight that the spread in the *ASPCAP* label distribution is strongly impacted by stars with modest S/N (< 120). The same result is not seen with *The Cannon*.

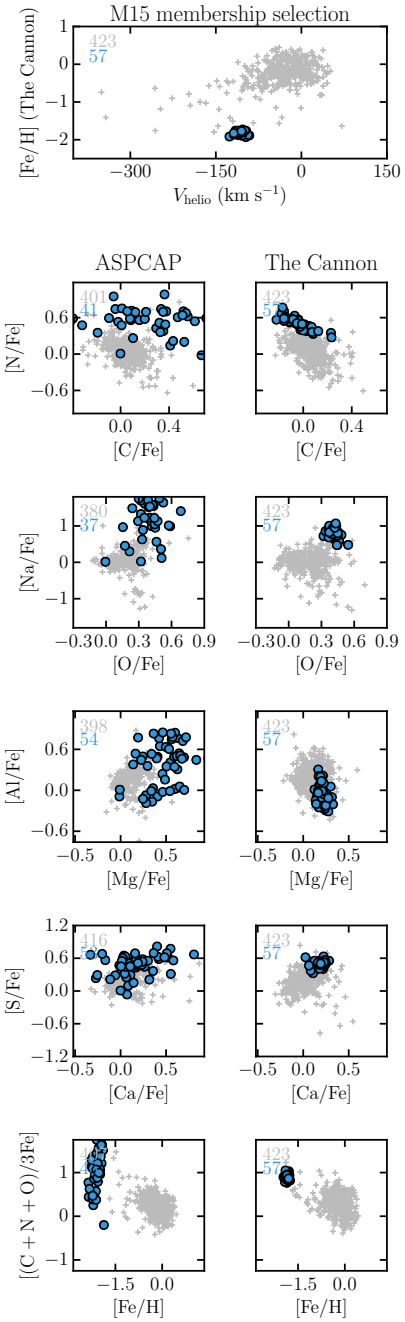

Fig. 12.— Abundance labels from *ASPCAP* and *The Cannon* for candidate members of the metal-poor globular cluster M 15 (NGC 7078). The top panel shows heliocentric velocity and metallicity, which we use to separate bonafide cluster members (blue) from field stars (grey). Panel limits are set by the dynamic range of *The Cannon* labels. The number of stars within those limits (from *ASPCAP* and *The Cannon*) are shown.

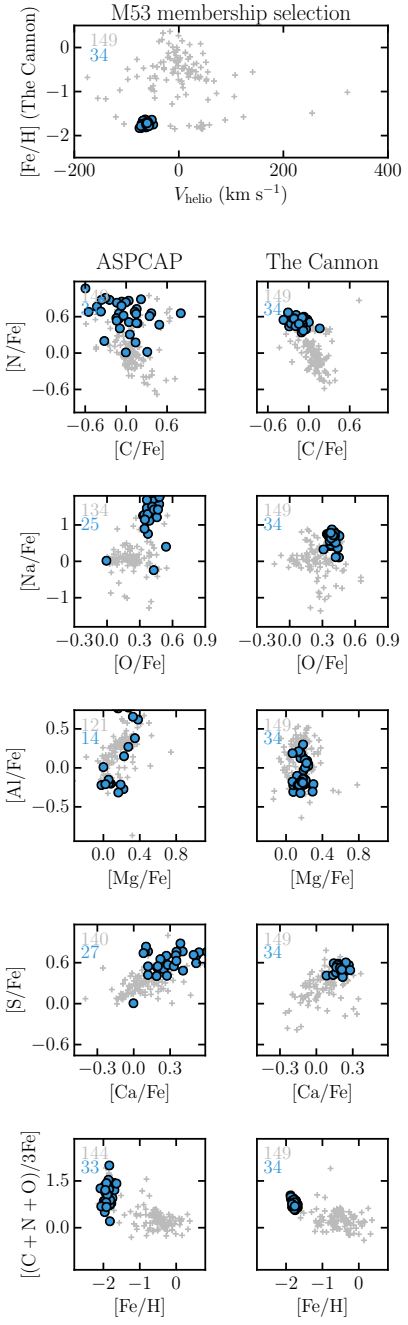

Fig. 13.— Label comparison between *ASPCAP* and *The Cannon* for M 53. Markings are the same as per Figure 12.

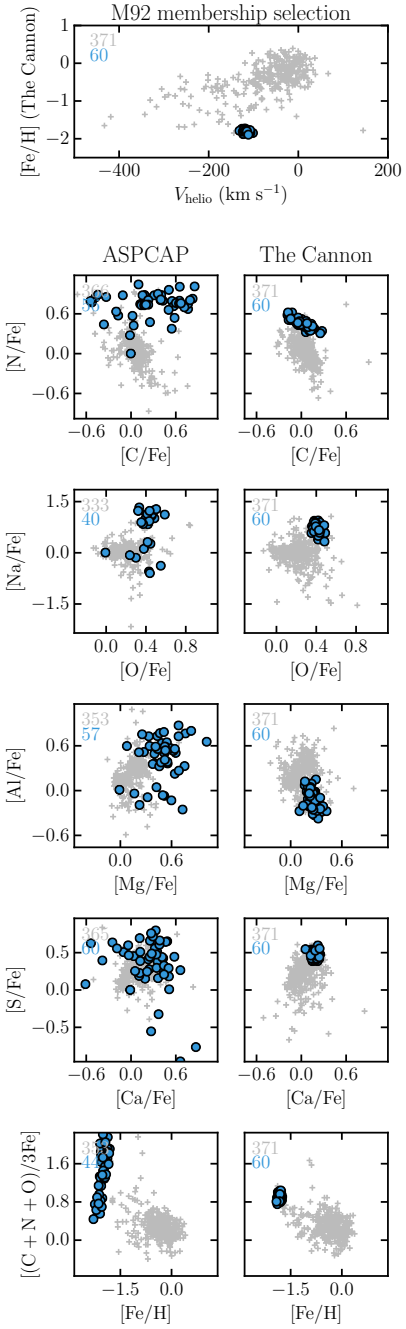

Fig. 14.— Label comparison between *ASPCAP* and *The Cannon* for M 92. Markings are the same as per Figure 12.

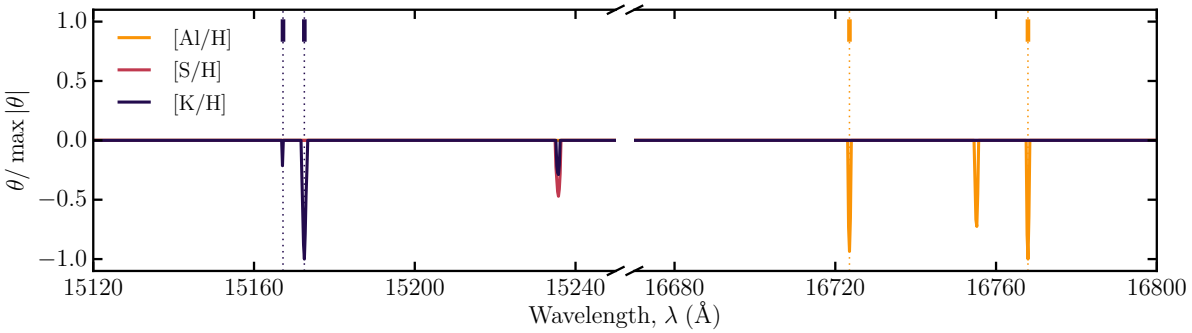

Fig. 15.— The normalized (to the maximum derivative value at any θ) first-order spectral derivatives for [Al/H], [S/H], and [K/H] from our regularized *Cannon* model. Vertical markings (and their colors) correspond to strong spectral lines used by Smith et al. (2013). For the sake of clarity we only show θ values above a background threshold.

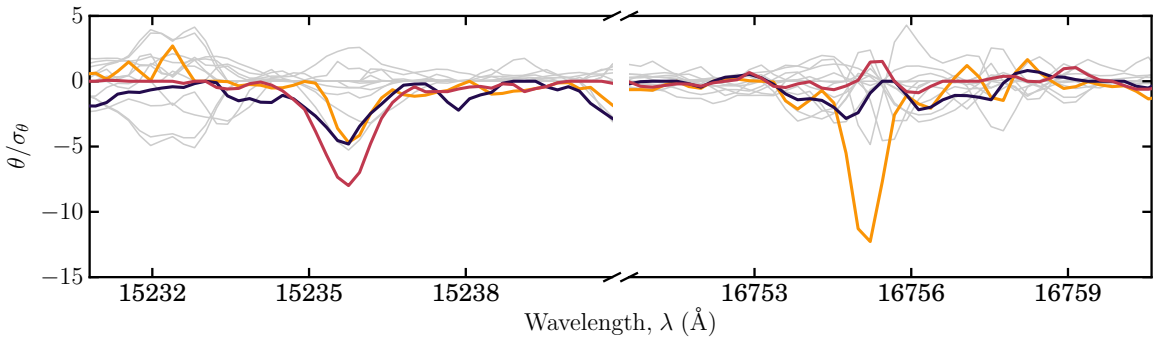

Fig. 16.— First-order spectral derivatives (normalized to the standard deviation for that derivative) for all abundance labels centered at the *APOGEE* 'unknown' lines at 15235.8 Å and 16755.6 Å. Abundance derivatives are shown in grey, or colored as per Figure 15.

Label	Error floor
$T_{\rm eff}$	22 K
$\log g$	0.03 dex
[C/H]	0.03 dex
[N/H]	0.04 dex
[O/H]	0.03 dex
[Na/H]	0.08 dex
[Mg/H]	0.03 dex
[Al/H]	0.06 dex
[Si/H]	0.03 dex
[S/H]	0.05 dex
[K/H]	0.07 dex
[Ca/H]	0.05 dex
[Ti/H]	0.07 dex
[V/H]	0.16 dex
[Mn/H]	0.04 dex
[Fe/H]	0.01 dex
[Ni/H]	0.03 dex

Table 1: Minimum error floor in labels, taken as the error in estimating labels for validation set stars at S/N = 50 (Figure 7).

APOGEE_ID	$T_{\rm eff}$ (K)	$\log g$	[Al/H]	[Ca/H]	···
2M18482322+4727598	4502 ± 3	2.279 ± 0.008	0.275 ± 0.013	-0.019 ± 0.012	···
2M18495146+4622483	4592 ± 5	2.641 ± 0.010	0.385 ± 0.016	0.109 ± 0.018	···
2M18543638+4658349	4767 ± 7	2.550 ± 0.014	-0.076 ± 0.019	-0.307 ± 0.017	···
2M18471131+4726561	4494 ± 3	2.158 ± 0.009	0.040 ± 0.014	-0.187 ± 0.013	···
2M18493578+4838112	4720 ± 6	2.828 ± 0.011	0.055 ± 0.019	-0.074 ± 0.017	···
2M18442845+4651265	4659 ± 5	2.559 ± 0.010	0.042 ± 0.014	-0.046 ± 0.013	···
2M18464037+4816449	4730 ± 4	2.621 ± 0.009	0.372 ± 0.013	0.204 ± 0.013	···
2M18515439+4809330	4786 ± 6	2.601 ± 0.013	0.009 ± 0.017	-0.208 ± 0.016	···

Table 2: All 17-labels for 87,563 *APOGEE* giant stars using the regularized model. Positions and ancillary measurements from *ASPCAP* are included for convenience. Only formal errors are listed here. See text for full details regarding the known error budget. This table is online-only. A portion is shown here to indicate form and content.